

1. Set-up

1. Mesh sizes

2. TODO

3. Runs that worked

1. mbcoupler test.14-jason.readpart large mesh

4. Runs that died

1. mbcoupler test.7-opt
- mbcoupler test.13-opt-1.8-readpart
- mbcoupler test.13-opt-1.8-bcastdel
- mbcoupler test.13-dbg-readpart
- mbcoupler test.13-dbg
- mbcoupler test.13.hdf1.8.readpart
- mbcoupler test.14-jason.readpart using big mesh pair

5. Romio bug

6. Bug #2

7. Bug #3

Set-up

Input files (with explanations and run scripts for intrepid): <http://www.mcs.anl.gov/~acaceres/couptest/>

Note: On intrepid, I'm not able to properly pass mbcoupler_test's "read options" through the command line. The queue system seems to eat the semi-colons. Workaround has just been to hardcode the read options in mbcoupler_test.cpp. Using:

```
readOpts = "PARALLEL=READ_PART;PARTITION=PARALLEL_PARTITION;PARTITION_DISTRIBUTE;PARALLEL_RESOLVE_SHAR
```

Executables and data files are in my home directory (/gpfs/home/acaceres).

Building with:

```
./configure --host=powerpc64-unknown-linux-gnu --enable-optimize --without-netcdf --with-hdf5=/soft/app
```

Was also using /soft/apps/hdf5-1.6.6/, might have been causing errors.

Note. Part of process is:

- Check out from svn on intrepid
- tar up and scp to login.mcs.anl.gov
- untar and run autoreconf on login.mcs
- retar and scp back to intrepid

variables:

- debug/optimized
- hdf 1.6 vs 1.8
- smp, dual or vn mode
- moab version

table: numprocs vs runtime/failure table: same as above, to be edited

input file how to run

issues:

- timeout with low procs
- hdf5 error
- core dump
-

Mesh sizes

File ptet_0.h5m:

type	count	total	minimum	average
Edge	91177	4.4e+04	0.28873	0.47854
Tri	60998	6e+03	0.058684	0.098831
Tet	767277	1.1e+04	0.0029522	0.014395
1D Side	4786656	2.5e+06	0.2118	0.52121
Vertex	134592			

File pvf-hex_0.h5m:

type	count	total	minimum	average
Edge	55332	2.6e+04	0.246	0.46534
Quad	27830	6e+03	0.081641	0.21658
Hex	110209	1.1e+04	0.037972	0.10022
1D Side	1433828	6.7e+05	0.246	0.46439
Vertex	117216			

TODO

add diagnostic to coupler.cpp to print out points being queried -- but why? we know the points are the vertices of the target mesh indep. program that loads a mesh in parallel and does point-query (from file or random points)

write debug output before and after communication. print out mpi communications done in mbcoupler_test.cpp (sanity check)

Runs that worked

Timings of successful runs of mbcoupler_test, organized by executable. First number in each row is instantiation time, others don't matter for now. Text on the right is "-n <number of nodes> -m <cores per node>" with smp=1, dual=2, vn=4.

==== mbcoupler_test.14-jason.readpart small mesh pair====

```
115.536  5.09885  0.196412  4 - vn
115.542  5.1021   0.193208  4 - vn
56.4102  1.92667  0.0981715 8 - vn
56.4118  1.92766  0.0979277 8 - vn
```

28.0366	0.979371	0.0722445	16	-vn
28.0389	0.98086	0.0722172	16	-vn
15.0396	0.692857	0.0420998	32	-vn
15.0408	0.691669	0.0415129	32	-vn
7.01743	0.434839	0.035589	64	-vn
7.01789	0.434329	0.036351	64	-vn
3.32072	0.33302	0.0305914	128	-vn
3.32112	0.336864	0.0339049	128	-vn
1.60273	0.240725	0.0424628	256	-vn
1.60459	0.23797	0.0436944	256	-vn
0.782535	0.139929	0.0354579	512	-vn
0.782575	0.146109	0.0426446	512	-vn
0.336923	0.12673	0.0494954	1024	-vn
0.338331	0.121506	0.0447295	1024	-vn
0.166049	0.0920526	0.0376375	2048	-vn
0.166185	0.0969077	0.0370007	2048	-vn
0.0892244	0.0964578	0.0353203	4096	-vn
0.0895257	0.100212	0.0439926	4096	-vn

mbcoupler_test.14-jason.readpart large mesh

1132.06	40.3868	1.7539	32
1132.06	40.4338	1.75465	32
1132.13	40.3718	1.75576	32
557.936	15.2882	1.04389	64
284.601	9.32169	0.609045	128
284.604	9.26964	0.611008	128
137.771	4.79621	0.339907	256
137.778	4.79567	0.339039	256
66.4713	2.31155	0.199686	512
66.4915	2.32137	0.202376	512
33.5435	1.42498	0.128231	1024
33.552	1.41945	0.122791	1024
15.5663	1.25793	0.0854221	2048
7.38878	0.713247	0.0632454	4096
7.39885	0.715164	0.0725397	4096
3.55494	0.671379	0.057236	8192
3.56014	0.669051	0.0538699	8192
1.96593	0.687914	0.0490007	16384

Runs that died

Unsuccessful runs, grouped by executable. Not clear yet if "maxtime exceeded" always means "the program was still thinking" versus "there was a deadlock". The smallest maxtime used was 10 minutes.

mbcoupler_test.7-opt

maxtime exceeded - read 0 mesh files -n 4 -m smp

mbcoupler_test.13-opt-1.8-readpart

maxtime exceeded - read 2 mesh files -n 1 -m smp
maxtime exceeded - read 2 mesh files -n 1 -m smp
maxtime exceeded - read 2 mesh files -n 1 -m smp
maxtime exceeded - read 2 mesh files -n 1 -m smp
romio bug -n 1024 -m dual
romio bug -n 1024 -m smp
romio bug -n 1024 -m smp
romio bug -n 128 -m dual
romio bug -n 256 -m dual
romio bug -n 256 -m dual
romio bug -n 512 -m dual
romio bug -n 512 -m dual
romio bug -n 512 -m smp
romio bug -n 512 -m smp
romio bug -n 512 -m smp

mbcoupler_test.13-opt-1.8-bcastdel

171-pvf-hex_0-ptet_0 -n 128 -m smp
180-pvf-hex_0-ptet_0 -n 128 -m smp
181-pvf-hex_0-ptet_0 -n 128 -m smp
maxtime exceeded - read 1 mesh files -n 128 -m smp
maxtime exceeded - read 1 mesh files -n 128 -m smp
maxtime exceeded - read 1 mesh files -n 2 -m smp
maxtime exceeded - read 1 mesh files -n 2 -m smp
maxtime exceeded - read 2 mesh files -n 1 -m smp
maxtime exceeded - read 2 mesh files -n 1 -m smp
maxtime exceeded - read 2 mesh files -n 1 -m smp
maxtime exceeded - read 2 mesh files -n 4 -m smp
maxtime exceeded - read 2 mesh files -n 4 -m smp

mbcoupler_test.13-dbg-readpart

```
maxtime exceeded - read 2 mesh files -n 1 -m smp  
maxtime exceeded - read 2 mesh files -n 2 -m smp
```

mbcoupler_test.13-dbg

```
killed by user -n 4 -m smp  
killed by user -n 8 -m smp  
maxtime exceeded - read 1 mesh files -n 2 -m smp  
maxtime exceeded - read 1 mesh files -n 4 -m smp  
maxtime exceeded - read 2 mesh files -n 1 -m smp
```

mbcoupler_test.13.hdf1.8.readpart

```
maxtime exceeded - read 2 mesh files -n 1 -m smp  
maxtime exceeded - read 2 mesh files -n 1 -m smp  
maxtime exceeded - read 2 mesh files -n 1 -m smp  
maxtime exceeded - read 2 mesh files -n 1 -m smp  
maxtime exceeded - read 2 mesh files -n 1 -m vn  
maxtime exceeded - read 2 mesh files -n 2 -m smp  
maxtime exceeded - read 2 mesh files -n 2 -m smp  
maxtime exceeded - read 2 mesh files -n 2 -m smp  
maxtime exceeded - read 2 mesh files -n 2 -m smp  
maxtime exceeded - read 2 mesh files -n 2 -m smp  
maxtime exceeded - read 2 mesh files -n 2 -m smp  
maxtime exceeded - read 2 mesh files -n 4 -m smp  
maxtime exceeded - read 2 mesh files -n 4 -m smp  
maxtime exceeded - read 2 mesh files -n 4 -m smp  
maxtime exceeded - read 2 mesh files -n 4 -m smp  
maxtime exceeded - read 2 mesh files -n 4 -m smp  
maxtime exceeded - read 2 mesh files -n 8 -m smp  
romio bug -n 128 -m vn  
romio bug -n 16 -m vn  
romio bug -n 32 -m vn  
romio bug -n 512 -m smp  
romio bug -n 64 -m vn  
romio bug -n 8 -m vn
```

mbcoupler_test.14-jason.readpart using big mesh pair

```
code 139 8 - vn  
queue system error 32 - vn
```

```

queue system error 64 - vn
queue system error 2048 - vn
queue system error 16384 - vn
romio 2           16 - vn
romio3            4 - vn
romio3            4 - vn
romio3            4 - vn
romio3            4 - vn
romio3            64 - vn

```

Romio bug

Stack trace

```

-----
+++ID TGID 706, Core 0, Thread 1 State 40000000, Sched: 48000000
0x016a7070
clone
?:0

-----
+++ID TGID 706, Core 0, Thread 5 State 00000000, Sched: 48000000 Running
0x016b56f4
raise
../nptl/sysdeps/unix/sysv/linux/raise.c:67

0x0167483c
abort
/bgssys/drivers/V1R4M2_200_2010-100508P/ppc/toolchain/gnu/glibc-2.4/stdlib/abort.c:73

0x014e3af4
MPID_Abort
/bghome/bgbuild/V1R4M2_200_2010-100508P/ppc/bgp/comm/lib/dev/mpich2/src/mpid/dcmfd/src/misc/mpid_abort

0x014b1edc
MPII_Abort
/bghome/bgbuild/V1R4M2_200_2010-100508P/ppc/bgp/comm/lib/dev/mpich2/src/mpi/init/abort.c:118

0x014ba168
ADIOI_Malloc_fn
/bghome/bgbuild/V1R4M2_200_2010-100508P/ppc/bgp/comm/lib/dev/mpich2/src/mpi/romio/adio/common/malloc.c:118

0x014cb534
ADIOI_R_Exchange_data_alltoallv
/bghome/bgbuild/V1R4M2_200_2010-100508P/ppc/bgp/comm/lib/dev/mpich2/src/mpi/romio/adio/ad_bgl/ad_bgl_r

0x014ccc78
ADIOI_Read_and_exch
/bghome/bgbuild/V1R4M2_200_2010-100508P/ppc/bgp/comm/lib/dev/mpich2/src/mpi/romio/adio/ad_bgl/ad_bgl_r

```

```
0x015137b4
MPIOI_File_read_all
/bghome/bgbuild/V1R4M2_200_2010-100508P/ppc/bgp/comm/lib/dev/mpich2/src/mpi/romio/mpi-io/read_all.c:107

0x014b67c4
MPII_File_read_at_all
/bghome/bgbuild/V1R4M2_200_2010-100508P/ppc/bgp/comm/lib/dev/mpich2/src/mpi/romio/mpi-io/read_atall.c:5

0x012b8700
H5FD_mpio_read
H5FDmpio.c:0

0x012ae45c
H5FD_read
?:0

0x012a3620
H5F_block_read
?:0

0x0128d92c
H5D_mpio_select_read
?:0

0x0129071c
H5D_final_collective_io
H5Dmpio.c:0

0x0128e26c
H5D_inter_collective_io
H5Dmpio.c:0

0x0128e11c
H5D_contig_collective_io
?:0

0x0127f27c
H5D_contig_read
H5Dio.c:0

0x0127d5f4
H5D_read
H5Dio.c:0

0x0127cd24
H5Dread
?:0

0x01181d30
mhdf_readwrite
src/util.c:0

0x01182c9c
mhdf_read_data
?:0

0x0117d82c
mhdf_readSetDataWithOpt
?:0
```

```

0x0115474c
_ZN4moab8ReadHDF520find_sets_containingEiiilRNS_5RangeE
?:0

0x0114f120
_ZN4moab8ReadHDF520find_sets_containingERNS_5RangeE
?:0

0x01147478
_ZN4moab8ReadHDF517load_file_partialEPKNS_11ReaderIface5IDTagEiiirkns_11FileOptionsE
?:0

0x01146728
_ZN4moab8ReadHDF591load_fileEPKcPKjRKNS_11FileOptionsEPKNS_11ReaderIface10SubsetListEPKPNS_7TagInfoE
?:0

0x010495f0
_ZN4moab4Core16serial_load_fileEPKcPKjRKNS_11FileOptionsEPKNS_11ReaderIface10SubsetListEPKPNS_7TagInfoE
?:0

0x011dfa44
_ZN4moab12ReadParallel9load_fileEPPKciPKjiRSsRSt6vectorIiSaIiEEbbSA_RKNS_11FileOptionsEPKNS_11ReaderIfa
?:0

0x011de6b8
_ZN4moab12ReadParallel9load_fileEPPKciPKjRKNS_11FileOptionsEPKNS_11ReaderIface10SubsetListEPKPNS_7TagI
?:0

0x01001a40
main
?:0

0x0166cdb8
generic_start_main
../csu/libc-start.c:231

0x0166d02c
__libc_start_main
../sysdeps/unix/sysv/linux/powerpc/libc-start.c:137

0xfffffffffc
??
?:0

-----
+++ID Tgid 706, Core 1, Thread 2 State 40000000, Sched: 080000000

0x016a7070
clone
?:0

```

Bug #2

When reading with

PARALLEL=BCAST_DELETE;PARTITION=PARALLEL_PARTITION;PARTITION_DISTRIBUTE;PARALLEL_RES

one of the processors get a copy of the entire target mesh (could also be keeping entire source mesh). Example with 4 procs (see last line):

```
point location: wanted 31316 got 30885 locally, 431 remote, missing 0  
point location: wanted 28794 got 27109 locally, 1685 remote, missing 0  
point location: wanted 31722 got 31715 locally, 7 remote, missing 0  
point location: wanted 134592 got 49339 locally, 85253 remote, missing 0
```

Pretty sure BCAST_DELETE was not doing this before.

Bug #3

Occurs with BCAST_DELETE and 128 procs, after first file gets read.

```
-----  
+++ID TGID 115, Core 0, Thread 1 State 40000000, Sched: 48000000  
  
0x016a7070  
H5HF_get_obj_len  
?:0  
  
-----  
+++ID TGID 115, Core 0, Thread 5 State 00000000, Sched: 48000000 Running  
  
0x010863f4  
_ZN4moab14AdaptiveKDTree10build_treeERKNS_5RangeERjPKNS0_8SettingsE  
/intrepid-fs0/users/acaceres/persistent/src/PM13-hdf1.8/src/AdaptiveKDTree.cpp:1367  
  
0x01050424  
_ZSt22__uninitialized_copy_aIN9__gnu_cxx17__normal_iteratorIPPN4moab7TagInfoEST6vectorIS4_SaIS4_EEEE9_  
/bgsys/drivers/V1R4M2_200_2010-100508P/ppc/gnu-linux/lib/gcc/powerpc-bgp-linux/4.1.2/../../../../include  
  
0x011a8bfc  
_ZSt6fill_nIPdiET_S1_T0_RKT1_  
/bgsys/drivers/V1R4M2_200_2010-100508P/ppc/gnu-linux/lib/gcc/powerpc-bgp-linux/4.1.2/../../../../include  
  
0x011a5c7c  
_ZN9__gnu_cxx13new_allocatorISt13_Rb_tree_nodeISt4pairIKN4moab19abaqus_element_typeEjEEE8allocateEjPKv  
/bgsys/drivers/V1R4M2_200_2010-100508P/ppc/gnu-linux/lib/gcc/powerpc-bgp-linux/4.1.2/../../../../include  
  
0x0119e4a0  
_ZN4moab9WriteUtil24get_entity_list_pointersENS_5Range14const_iteratorES2_PPKjNS_14WriteUtilIface14Ent_<...>  
/intrepid-fs0/users/acaceres/persistent/src/PM13-hdf1.8/src/WriteUtil.cpp:917  
  
0x0119be7c  
_ZNSaIPN4moab7TagInfoEEC1ERKS2_  
/bgsys/drivers/V1R4M2_200_2010-100508P/ppc/gnu-linux/lib/gcc/powerpc-bgp-linux/4.1.2/../../../../include  
  
0x0119afdb  
_ZNSaIiEC1ERKS_<...>  
/bgsys/drivers/V1R4M2_200_2010-100508P/ppc/gnu-linux/lib/gcc/powerpc-bgp-linux/4.1.2/../../../../include  
  
0x0119a9c4  
_ZSt8_DestroyIN9__gnu_cxx17__normal_iteratorIPiSt6vectorIiSaIiEEEEiEvT_S7_SaIT0_E  
/bgsys/drivers/V1R4M2_200_2010-100508P/ppc/gnu-linux/lib/gcc/powerpc-bgp-linux/4.1.2/../../../../include  
  
0x011dfede  
_ZNSt6vectorIdSaIdEEC1IiEET_S3_RKS0_<...>
```

```
/bgsys/drivers/V1R4M2_200_2010-100508P/ppc/gnu-linux/lib/gcc/powerpc-bgp-linux/4.1.2/../../../../../include  
0x011de6b8  
_ZN4moab10ReadABAQUS17read_element_listEjj  
/intrepid-fs0/users/acaceres/persistent/src/PM13-hdf1.8/src/io/ReadABAQUS.cpp:1303  
  
0x01001a40  
_ZN5SaIPN4moab7TagInfoEED2Ev  
/bgsys/drivers/V1R4M2_200_2010-100508P/ppc/gnu-linux/lib/gcc/powerpc-bgp-linux/4.1.2/../../../../../include  
  
0x0166cdb8  
H5T_vlen_disk_setnull  
H5Tvlen.c:0  
  
0x0166d02c  
H5T_vlen_reclaim_recurse  
H5Tvlen.c:0  
  
0xfffffffcc  
??  
??:0  
  
-----  
++ID TGID 115, Core 1, Thread 2 State 40000000, Sched: 08000000  
  
0x016a7070  
H5HF_get_obj_len  
??:0  
  
-----  
++ID TGID 115, Core 2, Thread 3 State 40000000, Sched: 08000000  
  
0x016a7070  
H5HF_get_obj_len  
??:0  
  
-----  
++ID TGID 115, Core 3, Thread 4 State 40000000, Sched: 08000000  
  
0x016a7070  
H5HF_get_obj_len  
??:0
```